

Hypothesis Testing for Parsimonious Gaussian Mixture Models

Antonio Punzo ^{*} Ryan P. Browne [†] Paul D. McNicholas [‡]

Abstract

Gaussian mixture models with eigen-decomposed covariance structures make up the most popular family of mixture models for clustering and classification, i.e., the Gaussian parsimonious clustering models (GPCM). Although the GPCM family has been used for almost 20 years, selecting the best member of the family in a given situation remains a troublesome problem. Likelihood ratio tests are developed to tackle this problems. These likelihood ratio tests use the heteroscedastic model under the alternative hypothesis but provide much more flexibility and real-world applicability than previous approaches that compare the homoscedastic Gaussian mixture versus the heteroscedastic one. Along the way, a novel maximum likelihood estimation procedure is developed for two members of the GPCM family. Simulations show that the χ^2 reference distribution gives reasonable approximation for the LR statistics only when the sample size is considerable and when the mixture components are well separated; accordingly, following Lo (2008), a parametric bootstrap is adopted. Furthermore, by generalizing the idea of Greselin and Punzo (2013) to the clustering context, a closed testing procedure, having the defined likelihood ratio tests as local tests, is introduced to assess a unique model in the general family. The advantages of this likelihood ratio testing procedure are illustrated via an application to the well-known Iris data set.

Key words: Parsimonious Gaussian Mixtures, Closed Testing Procedures, Eigen Decomposition, Homoscedasticity, Likelihood-Ratio Tests.

1 Introduction

The Gaussian mixture model (see Sect. 2.1) has been extensively considered as a powerful device for clustering by typically assuming that each mixture component represents a group (or cluster or class) in the original data (cf. Titterton *et al.*, 1985, Fraley and Raftery 1998, and McLachlan and Basford, 1988); however, merging can also

^{*}Department of Economics and Business, University of Catania, Corso Italia 55, 95129 Catania, Italy. Tel.: +39-095-7537640, e.mail: antonio.punzo@unict.it

[†]Department of Mathematics & Statistics, University of Guelph, Guelph, Ontario, Canada, N1G 2W1. Tel.: +1-519-8244120, ext. 53034, e.mail: rbrowne@uoguelph.ca

[‡]Department of Mathematics & Statistics, University of Guelph, Guelph, Ontario, Canada, N1G 2W1. Tel.: +1-519-8244120, ext. 53136, e.mail: paul.mcnicholas@uoguelph.ca

be considered to allow more than one component to represent a class (e.g., Hennig, 2010). Its popularity is largely attributable to its computational and theoretical convenience, as well as the speed with which it can be implemented for many data sets. Attention on Gaussian mixtures significantly increased since the work of Celeux and Govaert (1995), who proposed a family of fourteen Gaussian parsimonious clustering models (GPCMs) obtained by imposing some constraints on eigen-decomposed component covariance matrices. Popular software soon emerged for efficient implementation of some members of the GPCM family and served to further bolster their popularity (cf. Fraley and Raftery, 2002).

The GPCM family can be regarded as containing three subfamilies: the *spherical family* with two members that have spherical components, the *diagonal family* composed by four members that have axis-aligned components, and the *general family* with eight members that generate more flexible components. Homoscedasticity and heteroscedasticity represent the extreme configurations, in parsimonious terms, in the general family (see Sect. 2.1). Celeux and Govaert (1995) describe maximum likelihood (ML) estimation for the models in the general family (see Sect. 2.2); however, for two of these models, the authors relax one of the assumptions on which the family is based on, i.e., the assumption of decreasing order of the eigenvalues on the diagonal of the eigenvalues matrix. To overcome this problem, ML parameter estimation under order constraints is here proposed and illustrated (see Sect. 2.3).

When the number of components is either known *a priori* or determined by some of the methods available in the literature (see McLachlan and Peel, 2000, Chapt. 6 and the references therein), a likelihood-ratio (LR) statistic can be used for comparing the models in the general family. Unfortunately, attention has focused solely on the comparison between homoscedastic and heteroscedastic Gaussian mixtures and is further restricted to the univariate case (Lo, 2008). Herein, LR tests adopting the heteroscedastic Gaussian mixture model under the alternative hypothesis are considered for all the members of the GPCM family (Sect. 4). For these tests, simulation results show that the χ^2 reference distribution gives reasonable approximation for the LR statistic only when the sample size is considerable and when the mixture components are well separated (Sect. 3.1). This is expected within the mixture modelling context. In line with Lo (2008), a parametric bootstrap approach is so presented to approximate the distribution of the LR statistic (Sect. 3.2).

One drawback with the tests discussed above is that they are only pairwise tests, i.e., each model in the general family is separately compared with the benchmark heteroscedastic Gaussian mixture. An “overall” testing procedure that detects the model by simultaneously considering all the members of the general family is preferable. With this in mind, a closed testing procedure is developed based on the defined LR tests and building on recent work by Greselin and Punzo (2013) (Sect. 4). Computational aspects related to the implementation of the single LR test, and also to the implementation of the closed testing procedure, are given in Sect. 5. The well-known Iris data set is considered in Sect. 6 to illustrate the procedure and to demonstrate its advantages.

2 The GPCM Family

2.1 The family

The distribution of a p -variate random vector \mathbf{X} from a mixture of k Gaussian distributions is

$$f(\mathbf{x}; \boldsymbol{\vartheta}) = \sum_{j=1}^k \pi_j \phi(\mathbf{x}; \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j), \quad (1)$$

where $\pi_j > 0$ is the mixing proportion of the j th component, with $\sum_{j=1}^k \pi_j = 1$, $\phi(\mathbf{x}; \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)$ is the Gaussian density, with mean $\boldsymbol{\mu}_j$ and covariance matrix $\boldsymbol{\Sigma}_j$, and $\boldsymbol{\vartheta} = \{\pi_j, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j\}_{j=1}^k$.

The Gaussian mixture model (1) can be overparametrized because there are $p(p+1)/2$ free parameters for each $\boldsymbol{\Sigma}_j$. Banfield and Raftery (1993) introduce parsimony by considering the eigen-decomposition

$$\boldsymbol{\Sigma}_j = \lambda_j \boldsymbol{\Gamma}_j \boldsymbol{\Delta}_j \boldsymbol{\Gamma}_j', \quad (2)$$

for $j = 1, \dots, k$, where $\lambda_j = |\boldsymbol{\Sigma}_j|^{1/p}$, $\boldsymbol{\Delta}_j$ is the scaled ($|\boldsymbol{\Delta}_j| = 1$) diagonal matrix of the eigenvalues of $\boldsymbol{\Sigma}_j$ sorted in decreasing order, and $\boldsymbol{\Gamma}_j$ is a $p \times p$ orthogonal matrix whose columns are the normalized eigenvectors of $\boldsymbol{\Sigma}_j$, ordered according to their eigenvalues. Each element in the right side of (2) has a different geometric interpretation: λ_j determines the volume of the cluster, $\boldsymbol{\Delta}_j$ its shape, and $\boldsymbol{\Gamma}_j$ its orientation. Celeux and Govaert (1995) impose constraints on the elements on the right-hand side of (2) to give a family of 14 Gaussian parsimonious clustering models. These 14 models include very specific special cases, e.g., $\boldsymbol{\Gamma}_j = \mathbf{I}$ (identity matrix), and more general constraints, e.g., $\boldsymbol{\Gamma}_j = \boldsymbol{\Gamma}$.

Herein, we focus on the more general constraints. To this end, consider the triplet $(\lambda_j, \boldsymbol{\Delta}_j, \boldsymbol{\Gamma}_j)$ and allow its elements to be equal (E) or variable (V) across components. This leads to a ‘general family’ $\widetilde{\mathcal{M}}$ of eight models detailed in Table 1. With this convention, writing EEV means that we consider groups with equal volume, equal shape, and different orientation. Figure 1 exemplifies the models providing a graphical representation in the

Table 1: Models in the general family $\widetilde{\mathcal{M}}$ described by their covariance restrictions.

M	Volume	Shape	Orientation	$\boldsymbol{\Sigma}_j$	ML	Free covariance parameters
EEE	Equal	Equal	Equal	$\lambda \boldsymbol{\Gamma} \boldsymbol{\Delta} \boldsymbol{\Gamma}'$	CF	$p(p+1)/2$
VEE	Variable	Equal	Equal	$\lambda_j \boldsymbol{\Gamma} \boldsymbol{\Delta} \boldsymbol{\Gamma}'$	IP	$k + p - 1 + p(p-1)/2$
EVE	Equal	Variable	Equal	$\lambda \boldsymbol{\Gamma}_j \boldsymbol{\Delta} \boldsymbol{\Gamma}_j'$	IP	$1 + k(p-1) + p(p-1)/2$
EEV	Equal	Equal	Variable	$\lambda \boldsymbol{\Gamma} \boldsymbol{\Delta}_j \boldsymbol{\Gamma}'$	CF	$p + kp(p-1)/2$
VVE	Variable	Variable	Equal	$\lambda_j \boldsymbol{\Gamma}_j \boldsymbol{\Delta} \boldsymbol{\Gamma}_j'$	IP	$kp + p(p-1)/2$
VEV	Variable	Equal	Variable	$\lambda_j \boldsymbol{\Gamma} \boldsymbol{\Delta}_j \boldsymbol{\Gamma}'$	IP	$k + p - 1 + kp(p-1)/2$
EVV	Equal	Variable	Variable	$\lambda \boldsymbol{\Gamma}_j \boldsymbol{\Delta}_j \boldsymbol{\Gamma}_j'$	CF	$1 + k(p-1) + kp(p-1)/2$
VVV	Variable	Variable	Variable	$\lambda_j \boldsymbol{\Gamma}_j \boldsymbol{\Delta}_j \boldsymbol{\Gamma}_j'$	CF	$kp(p+1)/2$

case $p = k = 2$. For each model $M \in \widetilde{\mathcal{M}}$, the parameters in (1) can be denoted by $\boldsymbol{\vartheta}_M = \{\pi_j, \boldsymbol{\mu}_j, \lambda_j^M, \boldsymbol{\Delta}_j^M, \boldsymbol{\Gamma}_j^M\}_{j=1}^k$.

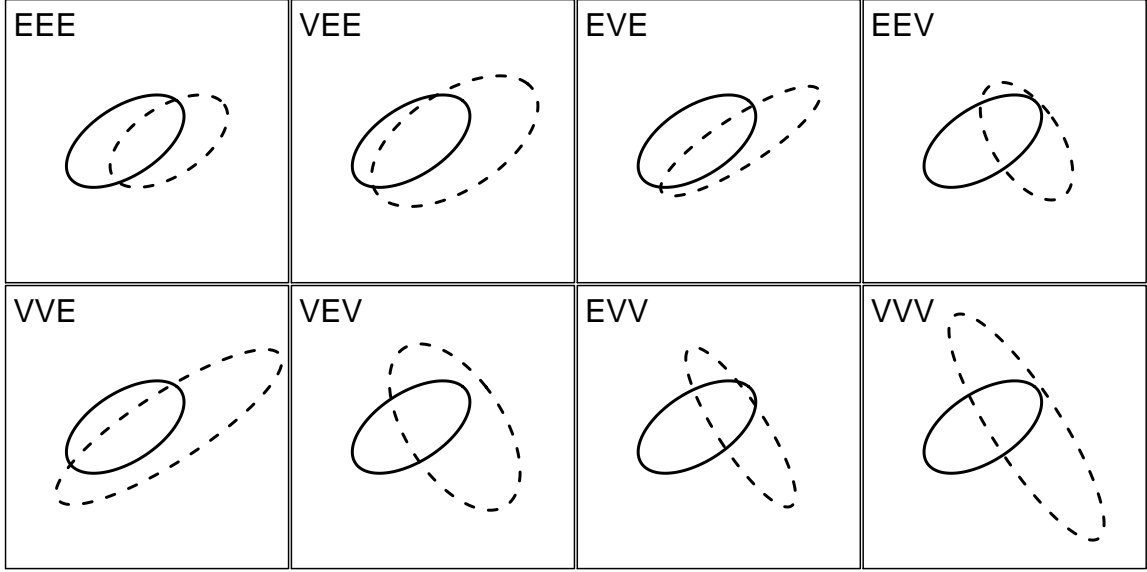


Figure 1: Examples of the models in $\widetilde{\mathcal{M}}$ in the bivariate case ($p = 2$) with $k = 2$ components.

2.2 Maximum likelihood parameter estimation

Given a sample $\mathbf{x}_1, \dots, \mathbf{x}_n$ from model (1), once k is assigned, the (observed-data) log-likelihood for the generic model $M \in \widetilde{\mathcal{M}}$ can be written as

$$l_M(\boldsymbol{\vartheta}_M) = \sum_{i=1}^n \log f(\mathbf{x}_i; \boldsymbol{\vartheta}_M). \quad (3)$$

The EM algorithm of Dempster *et al.* (1977) can be used to maximize l_M in order to find maximum likelihood (ML) estimates for $\boldsymbol{\vartheta}_M$. The algorithm basically works on the complete-data log-likelihood, i.e.,

$$l_M^c(\boldsymbol{\vartheta}_M) = \sum_{i=1}^n \sum_{j=1}^k z_{ij} [\log \pi_j + \log \phi(\mathbf{x}_i; \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j^M)], \quad (4)$$

where $z_{ij} = 1$ if \mathbf{x}_i comes from component j , and $z_{ij} = 0$ otherwise. The penultimate column of Table 3 indicates whether maximization, in the context of the generic M-step of the EM algorithm, can be achieved in a closed form (CF) or if an iterative procedure (IP) is needed; details can be found in Celeux and Govaert (1995) and Biernacki *et al.* (2008, pp. 22–24). However, for the models EVE and VVE, characterized by a common eigenvector matrix $\boldsymbol{\Gamma}$, Celeux and Govaert (1995) only describe an M-step for the weaker assumption of “equality in the set of p eigenvectors between groups” while, as stated in Sect. 2.1, we need equality in the ordered set of p eigenvectors between groups. This is also a fundamental requirement for the general family to be closed (see Sect. 4 for details). To motivate this problem, we consider the EEV model in Table 2 (see also Figure 2). This configuration erroneously corresponds to the EEE model in the “modified” general family of Celeux and Govaert (1995). For a further discussion about this issue see Greselin *et al.* (2011), Greselin and Punzo (2013), and Bagnato *et al.* (2014).

Table 2: Example of EEV model in the case of two groups in two dimensions.

	Volume	Shape	Orientation	Comment
Group 1	$\lambda_1 = 1$	$\Delta_1 = \begin{bmatrix} 4 & 0 \\ 0 & 1/4 \end{bmatrix}$	$\Gamma_1 = \begin{bmatrix} \sqrt{2}/2 & -\sqrt{2}/2 \\ \sqrt{2}/2 & \sqrt{2}/2 \end{bmatrix}$	Same ordering and values for the scaled eigenvalues in Δ_1 and Δ_2 but $\Gamma_1 \neq \Gamma_2$
Group 2	$\lambda_2 = 1$	$\Delta_2 = \begin{bmatrix} 4 & 0 \\ 0 & 1/4 \end{bmatrix}$	$\Gamma_2 = \begin{bmatrix} -\sqrt{2}/2 & \sqrt{2}/2 \\ \sqrt{2}/2 & \sqrt{2}/2 \end{bmatrix}$	

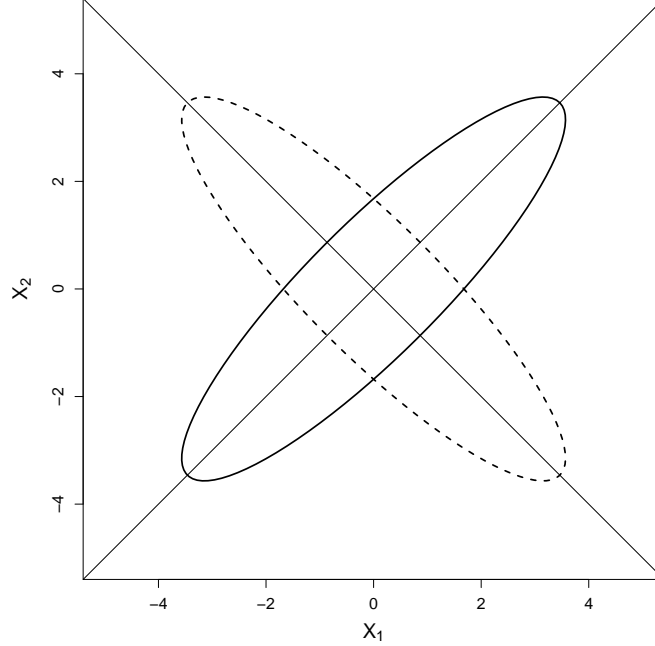


Figure 2: Ellipses related to the matrices in Table 2, where components have mean $(0, 0)'$.

2.3 Maximum likelihood parameter estimation under order constraints for EVE and VVE models

Motivated by the example given at the end of Sect. 2.2, we extend estimation procedures given in Celeux and Govaert (1995) for the models EVE and VVE to the case where we require the between-group eigenvalues to have the same decreasing order.

2.3.1 VVE Model: $\Sigma_j = \lambda_j \Gamma \Delta_j \Gamma'$

If we let $\Xi_j = \lambda_j \Delta_j$, where $\Xi_j = \text{diag}(\xi_{j1}, \dots, \xi_{jp})$, then maximizing the complete-data log-likelihood (4) is equivalent to minimizing the k optimization problems (one for each group j):

$$\begin{aligned} & \underset{\xi_{j1}, \dots, \xi_{jp}}{\text{minimize}} && \text{tr}(\mathbf{W}_j \Gamma \Xi_j^{-1} \Gamma') + n_j \log |\Xi_j| \\ & \text{subject to} && \xi_{j1} \geq \xi_{j2} \geq \dots \geq \xi_{jp} \geq 0, \end{aligned}$$

where, for the j th group, \mathbf{W}_j is the weighted scatter matrix and $n_j = \sum_{i=1}^n z_{ij}$, $j = 1, \dots, k$. This optimization problem is not convex because the second derivative of the

objective function can be negative. However, if we apply the transform $\zeta_{jl} = \log(\xi_{jl})$, which is one-to-one and ordering preserving, we obtain the convex programming problem:

$$\begin{aligned} & \underset{\zeta_{j1}, \dots, \zeta_{jp}}{\text{minimize}} && \sum_{l=1}^p [b_{jl} \exp(-\zeta_{jl}) + n_j \zeta_{jl}] \\ & \text{subject to} && \zeta_{j1} \geq \zeta_{j2} \geq \dots \geq \zeta_{jp}, \end{aligned} \quad (5)$$

where $b_{jl} = \boldsymbol{\gamma}_l' \mathbf{W}_j \boldsymbol{\gamma}_l$, with $\boldsymbol{\gamma}_l$ being the l th column vector of $\boldsymbol{\Gamma}$, also called the l th eigenvector. Advantageously, this convex programming problem has linear constraints and if a set of constraints are known to be active then the solution is easily to obtain. Thus, the primal active set method (Nocedal and Wright, 2000) is a good algorithm to apply this problem. Then to update the common orientation matrix, $\boldsymbol{\Gamma}$, one can apply the methodology from Flury and Gautschi (1986) and Browne and McNicholas (2014a,b).

2.3.2 EVE Model: $\boldsymbol{\Sigma}_j = \lambda \boldsymbol{\Gamma} \boldsymbol{\Delta}_j \boldsymbol{\Gamma}'$

For the case where the volume is equal across groups, we apply the same methodology from Sect. 2.3.1. However, in the transformed convex programming problem (5), for each j we have an additional constraint that $\sum_{l=1}^p \zeta_{jl} = 0$, which is equivalent to $\prod_{l=1}^p \xi_{jl} = 1$. This additional constraint is linear and can be adapted in the primal active set method.

3 Likelihood-ratio tests

Let $\overline{\mathcal{M}} = \widetilde{\mathcal{M}} / \{\text{VVV}\}$. For each $M \in \overline{\mathcal{M}}$, a natural way to test

$$H_0^M : \text{“data arise from } M\text{”} \quad \textit{versus} \quad H_1^{\text{VVV}} : \text{“data arise from VVV”},$$

consists of using the (generalized) likelihood-ratio (LR) statistic

$$\text{LR}_M = -2 \left[l_M(\widehat{\boldsymbol{\vartheta}}_M) - l_{\text{VVV}}(\widehat{\boldsymbol{\vartheta}}_{\text{VVV}}) \right], \quad (6)$$

where $\widehat{\boldsymbol{\vartheta}}_M$ and $\widehat{\boldsymbol{\vartheta}}_{\text{VVV}}$ are the ML estimators of $\boldsymbol{\vartheta}_M$ and $\boldsymbol{\vartheta}_{\text{VVV}}$ under the null and alternative hypotheses, respectively. Under some regularity conditions and under H_0^M , LR_M is commonly assumed asymptotically distributed as χ^2 with $\nu_M = \eta_{\text{VVV}} - \eta_M$ degrees of freedom, where η_{VVV} and η_M denote the number of (free) parameters for VVV and M , respectively. The value of ν_M is the gain in parsimony that could be achieved. Table 3 specifies the number of parameters η_M , and the degrees of freedom ν_M , for each $M \in \widetilde{\mathcal{M}}$.

3.1 Null distribution of the LR statistic

Unfortunately, with mixture models, regularity conditions may not hold for LR_M , $M \in \overline{\mathcal{M}}$, to have the assumed χ^2 reference distribution (Lo, 2008). Simulations are here conducted to examine this aspect. Because many factors come into play (e.g., the number of groups k , the dimension p of the observed variables, the overall sample size n , the

Table 3: Scheme of computation of ν_M , for the asymptotic χ^2 -approximation of LR_M , starting from η_{VVV} and η_M , $M \in \widetilde{\mathcal{M}}$.

M	η_{VVV}		η_M		ν_M
EEE	$(k-1) + kp + k\frac{p(p+1)}{2}$	–	$(k-1) + kp + \frac{p(p+1)}{2}$	=	$(k-1)\frac{p(p+1)}{2}$
VEE	$(k-1) + kp + k\frac{p(p+1)}{2}$	–	$(k-1) + kp + \frac{p(p+1)}{2} + (k-1)$	=	$(k-1)\left(\frac{p(p+1)}{2} - 1\right)$
EVE	$(k-1) + kp + k\frac{p(p+1)}{2}$	–	$(k-1) + kp + \frac{p(p+1)}{2} + (k-1)(p-1)$	=	$(k-1)\left(\frac{p(p-1)}{2} + 1\right)$
EEV	$(k-1) + kp + k\frac{p(p+1)}{2}$	–	$(k-1) + kp + k\frac{p(p+1)}{2} - (k-1)p$	=	$(k-1)p$
VVE	$(k-1) + kp + k\frac{p(p+1)}{2}$	–	$(k-1) + kp + \frac{p(p+1)}{2} + (k-1)p$	=	$(k-1)\frac{p(p-1)}{2}$
VEV	$(k-1) + kp + k\frac{p(p+1)}{2}$	–	$(k-1) + kp + k\frac{p(p+1)}{2} - (k-1)(p-1)$	=	$(k-1)(p-1)$
EVV	$(k-1) + kp + k\frac{p(p+1)}{2}$	–	$(k-1) + kp + k\frac{p(p+1)}{2} - (k-1)$	=	$(k-1)$
VVV	$(k-1) + kp + k\frac{p(p+1)}{2}$	–	$(k-1) + kp + k\frac{p(p+1)}{2}$	=	0

volume, shape, and orientation elements of the eigen-decomposition), some of them are necessarily considered fixed for our purposes.

One thousand data sets are generated from each model in $\overline{\mathcal{M}}$. We fix: $p = 2$, $k = 2$, $\pi_1 = \pi_2 = 0.5$, and $\boldsymbol{\mu}_1 = \mathbf{0}$. With regard to the remaining parameters of the models, in the bivariate case, we have

$$\boldsymbol{\Sigma}_j = \lambda_j \boldsymbol{\Gamma}_j \boldsymbol{\Delta}_j \boldsymbol{\Gamma}_j' = \lambda_j \mathbf{R}(\gamma_j) \begin{pmatrix} 1/\delta_j & 0 \\ 0 & \delta_j \end{pmatrix} \mathbf{R}(\gamma_j)',$$

where

$$\mathbf{R}(\gamma_j) = \begin{pmatrix} \cos \gamma_j & -\sin \gamma_j \\ \sin \gamma_j & \cos \gamma_j \end{pmatrix}$$

is the rotation matrix of angle γ_j , and $\delta_j \in (0, 1]$. Note that the elements in the shape matrix arise from the constraint $|\boldsymbol{\Delta}_j| = 1$. Hence, we have a single parameter for each element of the eigen-decomposition: λ_j is the volume parameter, δ_j is the shape parameter, and γ_j is the orientation parameter (for further details see Greselin *et al.*, 2011, and Greselin and Punzo, 2013). To generate data from each model, we preliminarily set $\boldsymbol{\Sigma}_1$ according to the values $\lambda_1 = 1$, $\delta_1 = 0.7$, and $\gamma_1 = \pi/6$ (i.e., 30°). With regard to $\boldsymbol{\Sigma}_2$, we choose $\lambda_2 = 3$ for models with variable volume, $\delta_2 = 0.3$ for models with variable shape, and $\gamma_2 = \pi/6 + \pi/4$ (i.e., $30^\circ + 45^\circ$) for models with variable orientation. The second variate μ_{22} of $\boldsymbol{\mu}_2 = (0, \mu_{22})'$ is computed, via a numerical procedure, to guarantee a fixed overlap between groups. Following Greselin and Punzo (2013), we adopted the normalized measure of overlap

$$B = \exp(-B^*),$$

which takes values between 0 (absence of overlap) and 1 (complete overlap), where

$$B^* = \frac{1}{8} \delta(\boldsymbol{\mu}_1, \boldsymbol{\mu}_2; \boldsymbol{\Sigma}^M) + \frac{1}{2} \log \left(\frac{|\boldsymbol{\Sigma}^M|}{\sqrt{|\boldsymbol{\Sigma}_1^M| + |\boldsymbol{\Sigma}_2^M|}} \right)$$

is the (positive) measure of overlap of Bhattacharyya (1943), with $\Sigma^M = (\Sigma_1^M + \Sigma_2^M) / 2$. In particular, we consider three scenarios: $B = 0.05$, $B = 0.25$, and $B = 0.45$. Three values for the sample size are also used: $n = 100$, $n = 200$, and $n = 500$. All nine combinations of the factors B and n are taken into account in the simulations.

Because the results are similar across models, Figure 3 shows the simulated distribution function (SDF) of the p -values (computed on 1000 replications) for EEE model only. The obtained results are arranged as a matrix of plots where n increases moving from

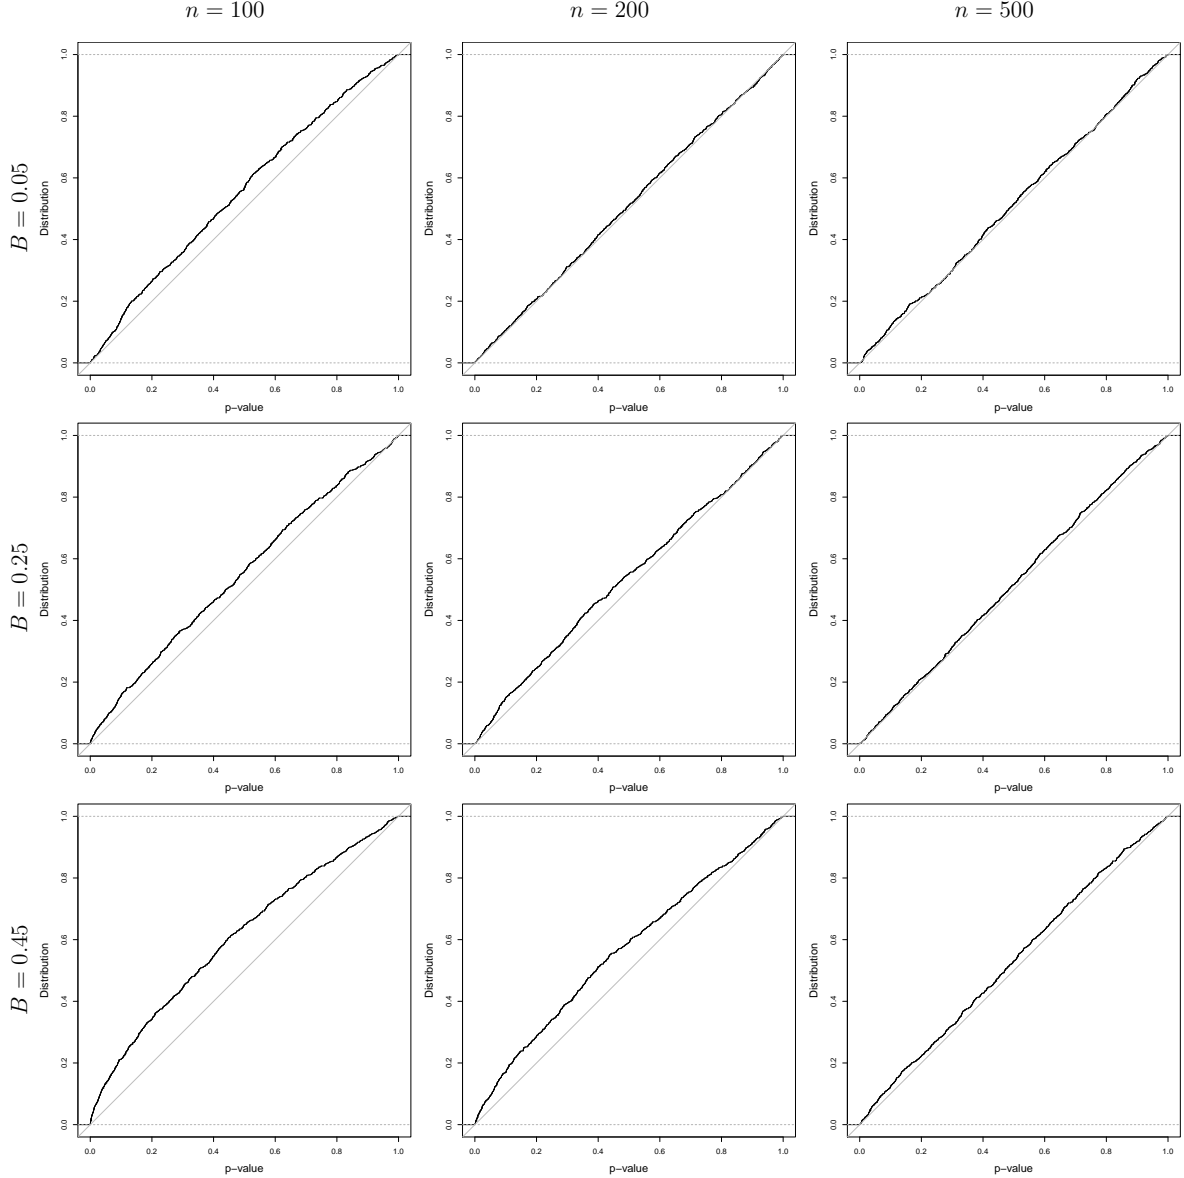


Figure 3: Asymptotic reference χ^2 distribution: simulated distribution function of the p -values, for model EEE, at the varying of sample size n and degree of overlap B .

left to right, while B increases moving from top to bottom. In each subplot, if the null distribution is well approximated by the χ^2 reference distribution, then we expect an SDF of the p -values very close to the distribution function of a uniform on $[0, 1]$, superimposed in gray in each subplot of Figure 3. These results suggest that the assumed χ^2 reference

distribution gives a good approximation when the sample size increases and/or when the degree of overlap decreases.

3.2 Parametric bootstrap LR tests

As the null χ^2 distribution does not provide a reasonable approximation for the LR statistics for small sample sizes (often encountered in practice) and for large overlap between groups, bootstrap methods that use the same criterion to compute LR_M for each bootstrap re-sample can be used to approximate the sampling distribution of LR_M .

In line with McLachlan (1987), McLachlan and Basford (1988, pp. 25–26), McLachlan and Peel (2000, Sect. 6.6) and Lo (2008), LR_M can be bootstrapped as follows. Proceeding under H_0^M , a bootstrap sample is generated from model M where $\boldsymbol{\vartheta}^M$ is replaced by its likelihood estimate formed under H_0^M from the original sample. The value of LR_M is computed, for the bootstrap sample, after fitting models M and VVV in turn to it. This process is repeated independently R times, and the replicated values of LR_M , formed from the successive bootstrap samples, provide an assessment of the null distribution of LR_M . This distribution enables an approximation to be made to the p -value corresponding to the value of LR_M evaluated from the original sample.

If a very accurate estimate of the p -value is required, then R should be large (Efron and Tibshirani, 1993). At the same time, when R is large, the amount of computation involved is considerable. However, there is usually no practical interest in estimating a p -value with high precision because the decision to be made concerns solely the rejection, or not, of H_0^M at a specified significance level α .

Aitkin *et al.* (1981) note that the bootstrap replications can be used to provide a test of approximate size α . In particular, the test that rejects H_0^M if LR_M for the original data is greater than the h th smallest of its R bootstrap replications has size

$$\alpha = 1 - \frac{h}{R+1}, \quad (7)$$

approximately. Hence, for a specified significance level α , the values of h and R can be chosen according to (7). For example, for $\alpha = 0.05$, the smallest value of R needed is 19 with $h = 19$. As cautioned above on the estimation of the p -value, R needs to be large to ensure an accurate assessment. In these terms, with $\alpha = 0.05$, the value $R = 99$ (and hence $h = 95$) could be a good compromise (McLachlan, 1987).

Under the same simulation design described in Sect. 3.1, Figure 4 shows the results of the parametric bootstrap approach with $R = 99$. Furthermore, in this case, because the obtained results are very similar across models, we report only those referred to model EEE. Figure 4 illustrates that the performance of the approach is always good regardless of the values of n and B .

4 Testing in the general family

So far we have discussed the assessment of each model in $\overline{\mathcal{M}}$ with respect to the (benchmark) alternative model VVV, which is the most general, unconstrained model. However, for real applications, we would prefer a statistical procedure to assess the true model with

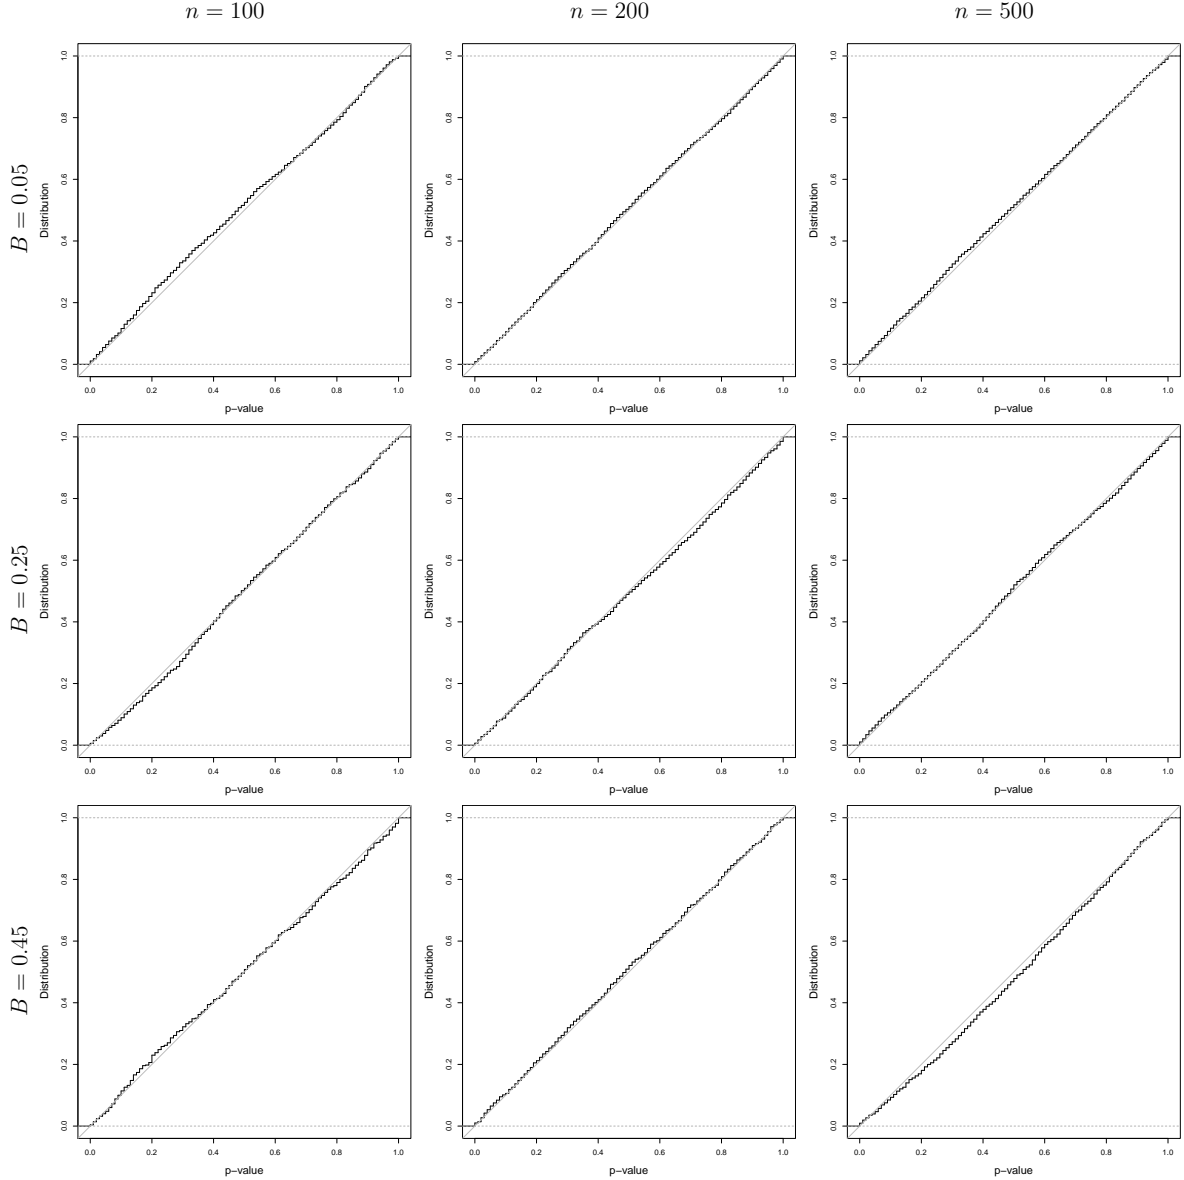


Figure 4: Parametric bootstrap: simulated distribution function of the p -values, for model EEE, at the varying of sample size n and degree of overlap B .

respect the entire general family $\widetilde{\mathcal{M}}$. To this end we generalize, to the clustering context, the closed LR testing procedure proposed by Greslin and Punzo (2013) for completely labeled data.

The hypotheses in

$$\mathcal{H} = \{H_0^{\text{VVE}}, H_0^{\text{VEV}}, H_0^{\text{EVV}}\}$$

are said to be *elementary* and, as detailed in Table 4, they play a crucial role: depending on the true model in $\widetilde{\mathcal{M}}$, none, some, or all of the hypotheses in \mathcal{H} may be the true null. Figure 5 also represents the null hypotheses H_0^M , $M \in \widetilde{\mathcal{M}}$, as a hierarchy where arrows indicate implications (see Hochberg and Tamhane, 1987, p. 344): for instance, H_0^{EEE} implies H_0^{VEE} , and this also means that model EEE is more restrictive (i.e., more parsimonious) than model VEE.

Table 4: Elementary hypotheses and their relationship with the models in $\widetilde{\mathcal{M}}$.

H_0^{EVV}	H_0^{VEV}	H_0^{VVE}		True model
True	True	True	\Rightarrow	EEE
False	True	True	\Rightarrow	VEE
True	False	True	\Rightarrow	EVE
True	True	False	\Rightarrow	EEV
False	False	True	\Rightarrow	VVE
False	True	False	\Rightarrow	VEV
True	False	False	\Rightarrow	EVV
False	False	False	\Rightarrow	VVV

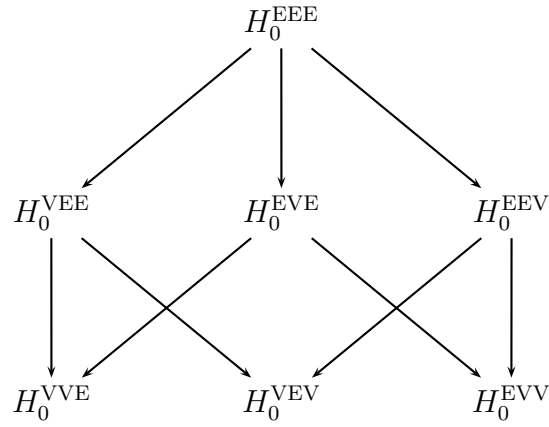


Figure 5: Graph of the hierarchy of relationships between the null hypotheses.

Operationally, according to the closed LR testing procedure of Greselin and Punzo (2013), we reject, say, the elementary hypothesis H_0^{VEV} if and only if each LR test on the more restrictive hypotheses H_0^{VEE} , H_0^{EEV} , H_0^{EEE} , and also on H_0^{VEV} itself, yields a significant result. Denoting by p_{VEE} , p_{EEV} , p_{EEE} , and p_{VEV} the p -values for H_0^{VEE} , H_0^{EEV} , H_0^{EEE} , and H_0^{VEV} , respectively, we report the *adjusted* p -value for H_0^{VEV} as $q_{\text{VEV}} = \max\{p_{\text{VEE}}, p_{\text{EEV}}, p_{\text{EEE}}, p_{\text{VEV}}\}$. An adjusted p -value represents the natural counterpart, in the multiple testing framework, of the classical p -value (see, e.g., Bretz *et al.*, 2011, p. 18). Specifically:

- they provide information about whether $H_0 \in \mathcal{H}$ is significant or not (q_M can be compared directly with any chosen significance level α and if $q_M \leq \alpha$, then H_0^M is rejected);
- they indicate “how significant” the result is (the smaller q_M , the stronger the evidence against H_0^M); and
- they are interpretable on the same scale as those for tests of individual hypotheses, making comparison with single hypothesis testing easier.

This closed testing procedure is the most powerful, among the available multiple testing procedures, that strongly controls the *familywise error rate* (FWER) at level α (as recently further corroborated via simulations by Giancristofaro Arboretti *et al.*, 2012). Controlling the FWER in a strong sense means controlling the probability of committing at least one Type I error under any partial configuration of true and false null hypotheses in \mathcal{H} . This is the only way to make inference on each hypothesis in \mathcal{H} . For further details on the closed testing procedure, and on its properties, see Greselin and Punzo (2013).

5 Computational aspects

Code for the LR tests (in both their χ^2 -based and bootstrap variants) and the closed LR testing procedure was written in the R computing environment (R Core Team, 2013). While specific code was written to obtain ML parameter estimation for models EVE and VVE (cf. Section 2.3), the **mixture** package (Browne and McNicholas, 2013) was used for the other models of the general family.

5.1 Initialization

5.1.1 LR tests

For model M , among the possible initialization strategies, each of the n vectors $\mathbf{z}_i^{(0)} = (z_{i1}^{(0)}, \dots, z_{ik}^{(0)})'$ can be randomly generated either in a “soft” way by generating k positive values summing to one, or in a “hard” way by randomly drawn a single observation from a multinomial distribution with probabilities $(1/k, \dots, 1/k)'$; see Biernacki *et al.* (2003), Karlis and Xekalaki (2003), and Bagnato and Punzo (2013) for more complicated strategies. Let $\hat{\mathbf{z}}_i^M$, $i = 1, \dots, n$, be the estimated posterior probabilities for model M . Because model M implies model VVV (that is M is nested in VVV), the “soft” values of $\hat{\mathbf{z}}_i^M$ can be used to initialize the EM algorithm for VVV; this forces, thanks to the monotonicity property of the EM algorithm (see, e.g., McLachlan and Krishnan, 2007), l_{VVV} to be greater than l_M and, hence, LR_M to be a well-defined positive value.

In the generic bootstrap re-sample from the fitted model M on the observed sample, we naturally know the true group membership of the generated observations. Thus, we can use the corresponding true “hard” values of \mathbf{z}_i , $i = 1, \dots, n$, to initialize the EM algorithm for model M . Once it is fitted, according to what said above, we can adopt the estimated posterior probabilities to initialize the EM algorithm for model VVV.

5.1.2 Closed testing procedure

With regard to the computation of the seven LR statistics in the closed testing procedure, on the observed sample we can take advantage of the hierarchy in Figure 5 to initialize the EM algorithms (for the use of hierarchical initialization strategies in mixture models see Ingrassia *et al.*, 2014 and Subedi *et al.*, 2013). In particular:

1. a “soft” or “hard” random initialization is used for model EEE in the top of the hierarchy;

2. the estimated posterior probabilities \hat{z}_i^{EEE} , $i = 1, \dots, n$, are used to initialize the EM algorithm for the models of the second level on the hierarchy (VEE, EVE, and EEV);
3. the posterior probabilities of the model with the highest log-likelihood between VEE and EVE are used to initialize the EM algorithm for VVE; the posterior probabilities of the model with the highest log-likelihood between VEE and EEV are used to initialize the EM algorithm for VEV; the posterior probabilities of the model with the highest log-likelihood between EVE and EEV are used to initialize the EM algorithm for EVV;
4. the posterior probabilities of the model with the highest log-likelihood between EVV, VEV, and VVE, are used to initialize the EM algorithm for VVV.

The described hierarchical initialization guarantees the natural ranking of log-likelihoods l_M , $M \in \mathcal{M}$.

5.2 Convergence criterion

The Aitken acceleration (Aitken, 1926) is used to estimate the asymptotic maximum of the log-likelihood at each iteration of the EM algorithm. Based on this estimate, we can decide whether or not the algorithm has reached convergence, i.e., whether or not the log-likelihood is sufficiently close to its estimated asymptotic value. For model $M \in \widetilde{\mathcal{M}}$, the Aitken acceleration at iteration $q + 1$, $q = 0, 1, \dots$, is given by

$$a_M^{(q+1)} = \frac{l_M^{(q+2)} - l_M^{(q+1)}}{l_M^{(q+1)} - l_M^{(q)}},$$

where $l_M^{(q+2)}$, $l_M^{(q+1)}$, and $l_M^{(q)}$ are the log-likelihood values from iterations $q + 2$, $q + 1$, and q , respectively. Then, the asymptotic estimate of the log-likelihood at iteration $q + 2$ is given by

$$l_{M,\infty}^{(q+2)} = l_M^{(q+1)} + \frac{1}{1 - a_M^{(q+1)}} \left(l_M^{(q+2)} - l_M^{(q+1)} \right),$$

cf. Böhning *et al.* (1994). The EM algorithm can be considered to have converged when $l_{M,\infty}^{(q+2)} - l_M^{(q+1)} < \epsilon$ (see Lindsay, 1995 and McNicholas *et al.*, 2010).

6 Analysis on the Iris data

In this section, we will show an application of the closed LR testing procedure on real data. A nominal level of 0.05 is adopted for the FWER-control and $R = 999$ bootstrap replications are considered; these values lead to $h = 950$ in (7). For completeness, the likelihood-based information criteria (IC) summarized in Table 5 will be also provided. In the definition of the ICL, $\text{MAP}(\hat{z}_{ij}^M) = 1$, if $\max_{h=1,\dots,k} \{\hat{z}_{ih}^M\}$ occurs at component j , and $\text{MAP}(\hat{z}_{ij}^M) = 0$ otherwise.

The Iris data set was made famous by Fisher (1936) as an illustration of discriminant analysis. Attention is here focused on the sample of $n = 100$ iris subdivided in $k = 2$

Table 5: Definitions and references for the adopted likelihood-based information criteria.

IC	Definition	Reference
AIC	$2l_M - 2\eta_M$	Akaike (1973)
AIC ₃	$2l_M - 3\eta_M$	Bozdogan (1994)
AICc	$AIC - 2\frac{\eta_M(\eta_M + 1)}{n - \eta_M - 1}$	Hurvich and Tsai (1989)
AICu	$AICc - n \log \frac{n}{n - \eta_M - 1}$	McQuarrie <i>et al.</i> (1997)
AWE	$2l_M - 2\eta_M \left(\frac{3}{2} + \log n \right)$	Banfield and Raftery (1993)
BIC	$2l_M - \eta_M \log n$	Schwarz (1978)
CAIC	$2l_M - \eta_M (1 + \log n)$	Bozdogan (1987)
ICL	$BIC + \sum_{i=1}^n \sum_{j=1}^k \text{MAP}(\hat{z}_{ij}^M) \log \hat{z}_{ij}^M$	Biernacki <i>et al.</i> (2000)

groups, of equal size, according to the species *versicolor* and *virginica*. On these flowers, $p = 4$ variables are measured in centimeters: *sepal length*, *sepal width*, *petal length*, and *petal width*. The matrix of scatter plots for these data is shown in Figure 6. A glance to this plot indicates that the groups are quite distinct, although some similarities in terms of orientation are observed. This conjecture is represented by model VVE.

We now work in the case $k = 2$ ignoring the known classification of the data. Table 6 contains the results from the closed LR testing procedure in both its χ^2 and bootstrap versions. First of all, we note that the null hypothesis H_0^{EEE} is rejected at any reasonable

Table 6: Details on the closed LR testing procedure applied to the Iris data. Bold numbers refer to the “not rejected” hypotheses in \mathcal{H} at the 0.05-level (columns q_M).

M	η_M	LR_M	ν_M	χ^2 -approximation		bootstrap	
				p_M	q_M	p_M	q_M
EEE	19	39.38134	10	0.00002		0.002	
VEE	20	24.97787	9	0.00300		0.012	
EVE	22	25.40289	7	0.00064		0.003	
EEV	25	26.05177	4	0.00003		0.001	
VVE	23	10.70523	6	0.09793	0.09793	0.155	0.155
VEV	26	11.89078	3	0.00777	0.00777	0.026	0.026
EVV	28	10.93548	1	0.00094	0.00094	0.010	0.010
VVV	29						

level ($p_{\text{EEE}} = 0.00002$ with the χ^2 -approximation, and $p_{\text{EEE}} = 0.002$ with the bootstrap approximation). Hence, if we limit the attention to this test only, as it is typically done in the literature, we should lean towards the adoption of a heteroscedastic Gaussian mixture

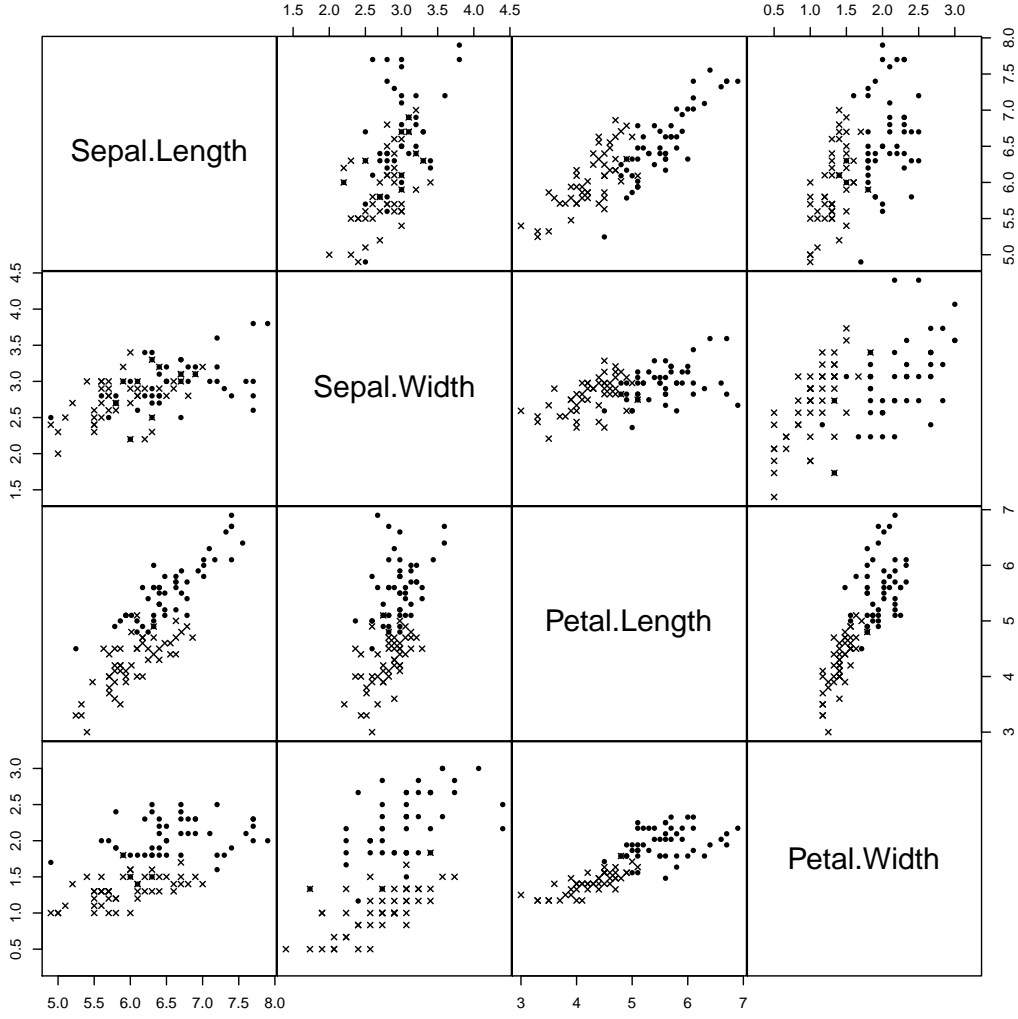


Figure 6: Matrix of scatter plots for a subset of Fisher's iris data (\times denotes *versicolor* and \bullet denotes *virginica*)

characterized by 29 parameters. On the contrary, additional information can be gained by looking at the closed testing procedure. In particular, Table 6 lists unadjusted p -values, for all the hypotheses in the hierarchy, and adjusted p -values for the elementary hypotheses. To facilitate the comprehension of how the adjusted p -values are computed, we can consider the following example in the bootstrap case: the adjusted p -value for H_0^{VEV} is given by

$$\begin{aligned} q_{\text{VEV}} &= \max\{p_{\text{VEV}}, p_{\text{EEV}}, p_{\text{VEE}}, p_{\text{EEE}}\} \\ &= \max\{0.026, 0.001, 0.012, 0.002\} = 0.026. \end{aligned}$$

At the 0.05-level, because H_0^{VEE} is the only elementary hypothesis that is not rejected in \mathcal{H} , it is also the hypothesis to be retained at the end of the procedure (see Table 4). This result does not vary by varying the approximation (χ^2 or bootstrap) of the LR statistics and it confirms our graphical conjectures. Moreover, it allows us to obtain a more parsimonious model (having 23 parameters) with a gain of 6 parameters with

respect to model VVV. Note also that, both the models (VVE and VVV) lead to five misallocated observations.

Table 7 reports the values of the likelihood-based information criteria of Table 5 for these data. Some concern arises when noting how different criteria can lead to different

Table 7: Likelihood-based information criteria for the Iris data. Bold numbers refer to the best model (highest column value) for each information criterion.

M	η_M	$2l_M$	AIC	AIC ₃	AICc	AICu	AWE	BIC	CAIC	ICL
EEE	19	-298.63	-336.63	-355.63	-346.13	-368.45	-530.63	-386.13	-405.13	-390.18
VEE	20	-284.23	-324.23	-344.23	-334.86	-358.43	-528.43	-376.33	-396.33	-378.62
EVE	22	-284.65	-328.65	-350.65	-341.80	-367.93	-553.28	-385.97	-407.97	-390.54
EEV	25	-285.30	-335.30	-360.30	-352.87	-382.98	-590.56	-400.43	-425.43	-404.13
VVE	23	-269.96	-315.96	-338.96	-330.48	-357.93	-550.79	-375.87	-398.87	-378.24
VEV	26	-271.14	-323.14	-349.14	-342.37	-373.84	-588.61	-390.88	-416.88	-392.80
EVV	28	-270.19	-326.19	-354.19	-349.06	-383.31	-612.07	-399.13	-427.13	-402.90
VVV	29	-259.25	-317.25	-346.25	-342.11	-377.77	-613.35	-392.80	-421.80	-394.40

choices; this consideration is further exacerbated if we consider that practitioners tend to use one of them almost randomly or routinely. On the other hand, the closed LR testing procedure offers a straightforward assessment of the model in the general family and it is based on only one subjective element, the significance level α , whose meaning is clear to everyone. Moreover, the adjusted p -values also provide a measure of “how significant” the test result is for each of the three terms of the eigen-decomposition: volume, shape, and orientation.

7 Discussion and future work

The likelihood-ratio statistic for comparing the homoscedastic Gaussian mixture *versus* its heteroscedastic version has been studied only in the univariate case (see Lo, 2008). Even if it were generalized to the multivariate case, being the resulting test *omnibus*, the practitioner should remain without any further information about a possible similarity across groups, different from the homoscedastic one, if the corresponding null hypothesis were rejected. Motivated by these considerations, in this paper we extended the use of likelihood-ratio tests in the multivariate case and to the general family of eight Gaussian mixture models of Celeux and Govaert (1995), homoscedasticity and heteroscedasticity being the extreme configurations in parsimony terms.

For two of the models in the general family, we also derived maximum likelihood parameter estimates fulfilling the requirements of the family — this is above and beyond the work of Celeux and Govaert (1995). For the resulting seven tests, we presented simulation results which were in line with those obtained by Lo (2008) in the univariate case for the likelihood-ratio test of homoscedasticity for Gaussian mixtures: the χ^2 reference distribution under the null does not provide a reasonable approximation for the likelihood-ratio statistic for small sample sizes and/or for large overlap between groups. To overcome this problem, we adopted a parametric bootstrap approach. Following work of Greselin and Punzo (2013) in the completely labeled scenario, the obtained tests were

also simultaneously considered in a closed testing procedure in order to assess a choice in the whole general family.

Although, in principle, an information criterion could be employed, a large number of these criteria have been proposed in literature (possibly leading to different choices as shown in the application to real data) and practitioners tend to use a given one of them routinely. On the other hand, the closed testing procedure illustrated here offers a straightforward assessment of the model in the general family and it is only based on one subjective element, the significance level α , whose meaning is clear to everyone. The real data set analyzed in the paper showed the gain in information and in parsimony that can be obtained by this approach.

A further remark refers to the type of application that is not restricted to model-based clustering. Our proposal provides indeed a suitable way to assess the model in the general family also for model-based classification — naturally based on Gaussian mixtures — where we fit our mixture models to data where some of the observations have known labels. In this case, we have also the advantage to know in advance the number of groups. Future work will investigate the extension of the closed testing procedure to the analogue general family for mixtures of t distributions (Andrews and McNicholas, 2012) and for mixtures of contaminated Gaussian distributions (Punzo and McNicholas, 2013).

References

- Aitken, A. (1926). On Bernoulli's numerical solution of algebraic equations. In *Proceedings of the Royal Society of Edinburgh*, volume 46, pages 289–305.
- Aitkin, M., Anderson, D., and Hinde, J. (1981). Statistical modelling of data on teaching styles. *Journal of the Royal Statistical Society. Series A (General)*, **144**(4), 419–461.
- Akaike, H. (1973). Information theory and an extension of maximum likelihood principle. In B. N. Petrov and F. Csaki, editors, *Second International Symposium on Information Theory*, pages 267–281, Budapest. Akademiai Kiado.
- Andrews, J. L. and McNicholas, P. D. (2012). Model-based clustering, classification, and discriminant analysis via mixtures of multivariate t -distributions. *Statistics and Computing*, **22**(5), 1021–1029.
- Bagnato, L. and Punzo, A. (2013). Finite mixtures of unimodal beta and gamma densities and the k -bumps algorithm. *Computational Statistics*, **28**(4), 1571–1597.
- Bagnato, L., Greselin, F., and Punzo, A. (2014). On the spectral decomposition in normal discriminant analysis. *Communications in Statistics - Simulation and Computation*, **43**(6), 1471–1489.
- Banfield, J. D. and Raftery, A. E. (1993). Model-based Gaussian and non-Gaussian clustering. *Biometrics*, **49**(3), 803–821.
- Bhattacharyya, A. (1943). On a measure of divergence between two statistical populations defined by their probability distributions. *Bulletin of the Calcutta Mathematical Society*, **35**(4), 99–109.

- Biernacki, C., Celeux, G., and Govaert, G. (2000). Assessing a mixture model for clustering with the integrated completed likelihood. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, **22**(7), 719–725.
- Biernacki, C., Celeux, G., and Govaert, G. (2003). Choosing starting values for the EM algorithm for getting the highest likelihood in multivariate Gaussian mixture models. *Computational Statistics & Data Analysis*, **41**(3-4), 561–575.
- Biernacki, C., Celeux, G., Govaert, G., Langrognet, F., Noulin, G., and Vernaz, Y. (2008). **MIXMOD** - *Statistical Documentation*. downloadable from http://www.mixmod.org/IMG/pdf/statdoc_2_1_1.pdf.
- Böhning, D., Dietz, E., Schaub, R., Schlattmann, P., and Lindsay, B. (1994). The distribution of the likelihood ratio for mixtures of densities from the one-parameter exponential family. *Annals of the Institute of Statistical Mathematics*, **46**(2), 373–388.
- Bozdogan, H. (1987). Model Selection and Akaike’s Information Criterion (AIC): The General Theory and its Analytical Extensions. *Psychometrika*, **52**, 345–370.
- Bozdogan, H. (1994). Mixture-model cluster analysis using model selection criteria and a new informational measure of complexity. In H. Bozdogan, editor, *Proceedings of the First US/Japan Conference on the Frontiers of Statistical Modeling: An Informational Approach*, pages 69–113, Netherlands. Springer-Verlag.
- Bretz, F., Hothorn, T., and Westfall, P. (2011). *Multiple Comparisons Using R*. Chapman & Hall, London.
- Browne, R. P. and McNicholas, P. D. (2013). ***mixture***: *Mixture Models for Clustering and Classification*.
- Browne, R. P. and McNicholas, P. D. (2014a). Estimating common principal components in high dimensions. *Advances in Data Analysis and Classification*, **8**(2).
- Browne, R. P. and McNicholas, P. D. (2014b). Orthogonal Stiefel manifold optimization for eigen-decomposed covariance parameter estimation in mixture models. *Statistics and Computing*, **24**(2), 203–210.
- Celeux, G. and Govaert, G. (1995). Gaussian parsimonious clustering models. *Pattern Recognition*, **28**(5), 781–793.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, **39**(1), 1–38.
- Efron, B. and Tibshirani, R. J. (1993). *An Introduction to the Bootstrap*, volume 57 of *Monographs on Statistics and Applied Probability*. Chapman & Hall, London.
- Fisher, R. A. (1936). The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, **7**(2), 179–188.

- Flury, B. N. and Gautschi, W. (1986). An algorithm for simultaneous orthogonal transformation of several positive definite matrices to nearly diagonal form. *SIAM Journal on Scientific and Statistical Computing*, **7**(1), 169–184.
- Fraley, C. and Raftery, A. E. (1998). How many clusters? Which clustering method? Answers via model-based cluster analysis. *Computer Journal*, **41**(8), 578–588.
- Fraley, C. and Raftery, A. E. (2002). Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association*, **97**(458), 611–631.
- Giancristofaro Arboretti, R., Bolzan, M., Bonnini, S., Corain, L., and Solmi, F. (2012). Advantages of the closed testing method in multiple comparisons procedures. *Communications in Statistics - Simulation and Computation*, **41**(6), 746–763.
- Greselin, F. and Punzo, A. (2013). Closed likelihood ratio testing procedures to assess similarity of covariance matrices. *The American Statistician*, **67**(3), 117–128.
- Greselin, F., Ingrassia, S., and Punzo, A. (2011). Assessing the pattern of covariance matrices via an augmentation multiple testing procedure. *Statistical Methods & Applications*, **20**(2), 141–170.
- Hennig, C. (2010). Methods for merging Gaussian mixture components. *Advances in Data Analysis and Classification*, **4**(1), 3–34.
- Hochberg, Y. and Tamhane, A. C. (1987). *Multiple Comparison Procedures*. Wiley, New York.
- Hurvich, C. M. and Tsai, C. L. (1989). Regression and time series model selection in small samples. *Biometrika*, **76**(2), 297–307.
- Ingrassia, S., Minotti, S. C., and Punzo, A. (2014). Model-based clustering via linear cluster-weighted models. *Computational Statistics & Data Analysis*, **71**, 159–182.
- Karlis, D. and Xekalaki, E. (2003). Choosing initial values for the EM algorithm for finite mixtures. *Computational Statistics & Data Analysis*, **41**(3–4), 577–590.
- Lindsay, B. (1995). *Mixture Models: Theory, Geometry and Applications*, volume 5. NSF-CBMS Regional Conference Series in Probability and Statistics, Institute of Mathematical Statistics, Hayward, California.
- Lo, Y. (2008). A likelihood ratio test of a homoscedastic normal mixture against a heteroscedastic normal mixture. *Statistics and Computing*, **18**(3), 233–240.
- McLachlan, G. J. (1987). On bootstrapping the likelihood ratio test statistic for the number of components in a normal mixture. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, **36**(3), 318–324.
- McLachlan, G. J. and Basford, K. E. (1988). *Mixture Models: Inference and Applications to Clustering*. Marcel Dekker, New York.

- McLachlan, G. J. and Krishnan, T. (2007). *The EM algorithm and extensions*. John Wiley & Sons, New York.
- McLachlan, G. J. and Peel, D. (2000). *Finite Mixture Models*. John Wiley & Sons, New York.
- McNicholas, P., Murphy, T., McDaid, A., and Frost, D. (2010). Serial and parallel implementations of model-based clustering via parsimonious Gaussian mixture models. *Computational Statistics & Data Analysis*, **54**(3), 711–723.
- McQuarrie, A., Shumway, R., and Tsai, C.-L. (1997). The model selection criterion AIC_c. *Statistics & Probability Letters*, **34**(3), 285–292.
- Nocedal, J. and Wright, S. J. (2000). *Numerical Optimization*. Springer.
- Punzo, A. and McNicholas, P. D. (2013). Robust clustering via parsimonious mixtures of contaminated Gaussian distributions. arXiv.org e-print 1305.4669, available at: <http://arxiv.org/abs/1305.4669>.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, **6**(2), 461–464.
- R Core Team (2013). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Subedi, S., Punzo, A., Ingrassia, S., and McNicholas, P. D. (2013). Clustering and classification via cluster-weighted factor analyzers. *Advances in Data Analysis and Classification*, **7**(1), 5–40.
- Titterton, D. M., Smith, A. F. M., and Makov, U. E. (1985). *Statistical Analysis of Finite Mixture Distributions*. John Wiley & Sons, New York.